

BONNSAI: correlated stellar observables in Bayesian methods

F.R.N. Schneider¹*, N. Castro², L. Fossati³, N. Langer², and A. de Koter^{4,5}

¹ Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, United Kingdom

² Argelander-Institut für Astronomie der Universität Bonn, Auf dem Hügel 71, 53121 Bonn, Germany

³ Space Research Institute, Austrian Academy of Sciences, Schmiedlstrasse 6, 8042 Graz, Austria

⁴ Astronomical Institute 'Anton Pannekoek', Amsterdam University, Science Park 904, 1098 XH, Amsterdam, The Netherlands

⁵ Instituut voor Sterrenkunde, KU Leuven, Celestijnenlaan 200D, 3001, Leuven, Belgium

Received day month 2016 / Accepted day month 2016

ABSTRACT

In an era of large spectroscopic surveys of stars and big data, sophisticated statistical methods become more and more important in order to infer fundamental stellar parameters such as mass and age. Bayesian techniques are powerful methods because they can match all available observables simultaneously to stellar models while taking prior knowledge properly into account. However, in most cases it is assumed that observables are uncorrelated which is generally not the case. Here, we include correlations in the Bayesian code BONNSAI by incorporating the covariance matrix in the likelihood function. We derive a parametrisation of the covariance matrix that, in addition to classical uncertainties, only requires the specification of a correlation parameter that describes how observables covary. Our correlation parameter depends purely on the method with which observables have been determined and can be analytically derived in some cases. This approach therefore has the advantage that correlations can be accounted for even if information for them are not available in specific cases but are known in general. Because the new likelihood model is a better approximation of the data, the reliability and robustness of the inferred parameters are improved. We find that neglecting correlations biases the most likely values of inferred stellar parameters and affects the precision with which these parameters can be determined. The importance of these biases depends on the strength of the correlations and the uncertainties. For example, we apply our technique to massive OB stars, but emphasise that it is valid for any type of stars. For effective temperatures and surface gravities determined from atmosphere modelling, we find that masses can be underestimated on average by 0.5σ and mass uncertainties overestimated by a factor of about 2 when neglecting correlations. At the same time, the age precisions are underestimated over a wide range of stellar parameters. We conclude that accounting for correlations is essential in order to derive reliable stellar parameters including robust uncertainties and will be vital when entering an era of precision stellar astrophysics thanks to the Gaia satellite.

Key words. Methods: data analysis – Methods: statistical – Stars: general – Stars: fundamental parameters

1. Introduction

With the advent of large stellar spectroscopic surveys, powerful telescopes, and advanced spectral modelling capabilities, more and more is known about individual stars with ever increasing accuracy and precision. In particular, the Gaia satellite will revolutionise the precision with which distances, luminosities, and other stellar parameters can be determined.

Accuracy and precision are essential to many astrophysical applications, for example when determining fundamental properties of exoplanets and the architectures of planetary systems from inferred masses, radii, and ages of their host stars (e.g. Bonfanti et al. 2016); when studying the dynamical and chemical evolution of the Galaxy from F and G stars (e.g. Mitschang et al. 2014); when testing stellar models with in-depth observations of binary stars (e.g. Lastennet & Valls-Gabaud 2002; Torres et al. 2010); or when investigating the consistency of different stellar age estimate methods (e.g. Maxted et al. 2015b). Also, in statistical studies of large samples of stars, systematic biases have the potential to hamper our interpretation of the data (e.g. Jørgensen & Lindegren 2005; Fossati et al. 2016).

In order to determine robust and reliable stellar parameters, sophisticated statistical methods are indispensable and systematic uncertainties need to be understood. This includes systematics in the stellar models (e.g. Martins & Palacios 2013; Jones

et al. 2015; Stancliffe et al. 2015, 2016), but also systematics in the statistical techniques used to compare observations of stars with stellar models. Bayesian statistical methods (e.g. Jørgensen & Lindegren 2005; da Silva et al. 2006; Takeda et al. 2007; Shkedy et al. 2007; van Dyk et al. 2009; Burnett & Binney 2010; Serenelli et al. 2013; Schönrich & Bergemann 2014; Schneider et al. 2014; Maxted et al. 2015a) have proven to be powerful because they can compare all available observables including error bars simultaneously to stellar models; they also take prior knowledge such as initial mass functions properly into account. Moreover, such methods have the potential to deliver robust and reliable error bars of inferred stellar parameters that do not suffer from biases, for example due to neglecting the time spent by stars in various regions of the Hertzsprung–Russell (HR) diagram (e.g. Pont & Eyer 2004).

For simplicity, it is in most cases assumed that observables are independent of each other, i.e. that they are uncorrelated. However, this assumption is not generally true and many observables are in fact correlated. For example, effective temperatures and surface gravities are correlated if inferred from stellar atmosphere modelling because they can influence diagnostic lines in a similar way, i.e. effective temperature and surface gravity depend on each other. The same is true when converting observed effective temperatures, distances, and magnitudes of stars into bolometric luminosities. To accomplish this it is necessary to use the bolometric correction, which is a function of effective tem-

* fabian.schneider@physics.ox.ac.uk

perature, because of the approximate black-body behaviour of stars. Therefore, the luminosity also depends on effective temperature: bolometric luminosity and effective temperature co-vary, i.e. they are correlated.

Correlations are valuable information and, when not accounted for, can bias the comparison of observations with stellar models. Correlations change the likelihood function and, as we show in this paper, thereby the precision and most likely values of inferred fundamental stellar parameters such as mass and age. They provide additional information on the observables that ultimately result in more reliable stellar parameters. Taking correlations properly into account may decrease the error bars such that stellar parameters can be inferred with higher precision.

BONNSAI¹ (Schneider et al. 2014) is a Bayesian statistical method that matches all the available observables of stars simultaneously to models of stellar evolution in order to determine fundamental stellar parameters such as mass and age, to predict yet unobserved stellar parameters, and to probe theses models using sophisticated goodness-of-fit tests. Being a Bayesian method, it takes prior knowledge such as initial mass functions properly into account. In this paper, we generalise the likelihood function of BONNSAI by adding the covariance matrix to be able to account for correlated stellar observables. So far, correlations of stellar parameters are typically not reported in the literature, but only marginalised 1σ uncertainties. In order to facilitate the use of such stellar parameters, we derive a parametrisation of the covariance matrix that depends on conventional 1σ uncertainties, and also on a correlation parameter that describes how two observables co-vary and that purely depends on the method with which the observables have been determined. Once the correlation parameter for a particular method is known, this parametrisation has the advantage that correlations can be accounted for when matching observables against stellar models even if only marginalised 1σ error bars are available. The new likelihood model can be applied in various situations (also outside stellar astrophysics) and is valid for any kind of star in any evolutionary state.

In Sec. 2 we describe the new likelihood model and compare it to the old one and to even more accurate numerical models. We investigate how the new likelihood model affects the precision and most likely values of inferred initial masses and ages in Sec. 3, and summarise our conclusions in Sec. 4. A version of BONNSAI that includes the new likelihood model is available through a web-interface².

2. Method

2.1. The new likelihood model

The heart of any Bayesian method such as BONNSAI is Bayes' theorem,

$$p(\mathbf{m}|\mathbf{d}) \propto p(\mathbf{d}|\mathbf{m})p(\mathbf{m}). \quad (1)$$

Bayes' theorem states that the posterior probability distribution $p(\mathbf{m}|\mathbf{d})$, i.e. the probability distribution of the model parameters \mathbf{m} given observational data \mathbf{d} , follows from the likelihood $p(\mathbf{d}|\mathbf{m})$, i.e. the probability distribution of the observational data given the model parameters and the prior distribution $p(\mathbf{m})$ that encompasses a priori knowledge about the model parameters.

In BONNSAI and similar methods it is usually assumed that the likelihood function can be written as a product of Gaussian

functions for each observable. This assumption is true if the observables are normally distributed and uncorrelated.

If the observables are correlated, the product of independent Gaussian likelihoods is only a zeroth order approximation of the data; this approximation may still be adequate to describe the data if the correlations are weak. In a first-order approximation, it is possible to generalise the Gaussian likelihood function by introducing the covariance matrix Σ ,

$$p(\mathbf{d}|\mathbf{m}) = \frac{1}{\sqrt{(2\pi)^k |\Sigma|}} \exp \left[-\frac{1}{2} (\mathbf{d} - \mathbf{d}(\mathbf{m}))^T \Sigma^{-1} (\mathbf{d} - \mathbf{d}(\mathbf{m})) \right], \quad (2)$$

where k is the dimension, i.e. the number of observables, $|\Sigma|$ the determinant of the covariance matrix, and $\mathbf{d}(\mathbf{m})$ the predicted observables for model parameters \mathbf{m} . If the observables are uncorrelated, the covariance matrix has only diagonal elements and the likelihood function in Eq. 2 reduces to a product of independent Gaussian functions. The likelihood model in Eq. 2 may still not be good enough to properly approximate the data, for example if the shape of the likelihood is that of a banana. In that case it is probably easiest to use a numerical model of the likelihood function in Bayes' theorem.

Covariances of inferred stellar parameters are typically not reported in the literature, but only 1σ uncertainties of individual observables; we call these 1σ uncertainties “conventional uncertainties/error bars”. Our aim is therefore to derive a parametrisation of the covariance matrix that is fully specified by conventional error bars and additional parameters that we call “correlation parameters” that only depend on the method with which observables have been derived. Once the correlation parameters are known for a given method, this approach will enable us to incorporate correlations even in such cases where only conventional error bars are available.

Covariances are often expressed in terms of Pearson's correlation parameter $\rho_{pq} = \text{Cov}(p, q) / \sigma_p \sigma_q$ where Cov is the covariance of observables p and q , and σ_p and σ_q are the respective conventional 1σ uncertainties. With this notation, the covariance matrix Σ has the square of the conventional uncertainties on its diagonal and $\rho_{pq}\sigma_p\sigma_q$ on its off-diagonal (recall that Σ is symmetric). In this way, the covariance matrix would be given by conventional uncertainties and Pearson's correlation parameter as desired; however, Pearson's correlation parameter depends not only on the method with which observables have been determined, but also on the conventional uncertainties such that this correlation parameter cannot be re-used in cases where only conventional uncertainties are available (see Sec. 2.1.1).

We therefore seek to derive an alternative representation of the covariance matrix for which correlation parameters can be re-used or derived analytically (see Sec. 2.3). To that end, we note that the two-dimensional version of the likelihood function defined in Eq. 2 has the shape of a rotated ellipse (Fig. 1). The semi-major and -minor axes a and b of this Gaussian ellipse define the orientation of a reference frame that is rotated with respect to the standard reference frame. Let \mathcal{B} be the standard Euclidean basis and \mathcal{T} the basis of the rotated reference frame (the basis \mathcal{T} is given by the eigenvectors of Σ and the eigenvalues are a^2 and b^2). In the rotated reference frame, the covariance matrix is diagonal and Eq. (2) can be written as a product of independent Gaussians whose variances are the eigenvalues of the covariance matrix.

There is a linear map $\mathbf{R} : \mathcal{B} \rightarrow \mathcal{T}$ that transforms from basis \mathcal{B} to \mathcal{T} and that can be expressed as a rotation matrix. Using this transformation \mathbf{R} we can change the basis of the inverse of the covariance matrix to switch between representations relative to

¹ The BONNSAI web-service is available at www.astro.uni-bonn.de/stars/bonnsai.

² www.astro.uni-bonn.de/stars/bonnsai

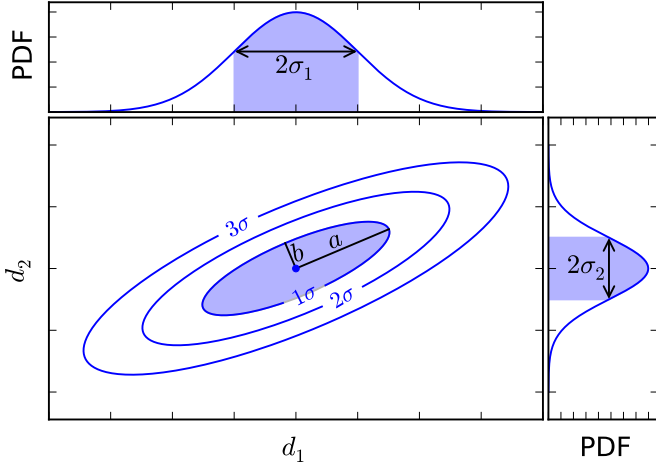


Fig. 1. Schematic representation of the 1σ , 2σ , and 3σ contours of a Gaussian likelihood function of two correlated observables d_1 and d_2 . The semi-major and -minor axes a and b of the 1σ contour are indicated. The marginalised probability density functions (PDFs) of the two observables are shown in the top and right panels. They are Gaussian functions with standard deviations σ_1 and σ_2 (the conventional error bars). In this example, $\sigma_1 = 2\sigma_2$.

the basis \mathcal{B} and \mathcal{T} ,

$$\Sigma_{\mathcal{B}}^{-1} = \mathbf{R} \Sigma_{\mathcal{T}}^{-1} \mathbf{R}^T, \quad (3)$$

where $\Sigma_{\mathcal{B}}$ and $\Sigma_{\mathcal{T}}$ denote the covariance matrix with respect to basis \mathcal{B} and \mathcal{T} , respectively. In most applications the conventional uncertainties σ of each observable are known but not the standard deviations of the likelihood in the rotated reference frame (σ_1 and σ_2 vs. a and b ; cf. Fig. 1). To compute one from the other, we derive relations between the correlation parameters, the conventional error bars, and the standard deviations of the rotated Gaussian likelihood by integrating the likelihood function in Eq. (2). In Sec. 2.1.1 we derive the expressions for two correlated observables and in Sec. 2.1.2 for three.

2.1.1. Two correlated observables

For two observables the rotation matrix \mathbf{R} can be written as

$$\mathbf{R} = \begin{pmatrix} \cos \varphi & -\sin \varphi \\ \sin \varphi & \cos \varphi \end{pmatrix} \quad (4)$$

with φ being the rotation angle by which the reference frame \mathcal{T} is rotated with respect to \mathcal{B} ; we call φ the correlation parameter. Let a and b be the standard deviations of the rotated Gaussian such that the inverse of the covariance matrix with respect to basis \mathcal{T} has the form

$$\Sigma_{\mathcal{T}}^{-1} = \begin{pmatrix} 1/a^2 & 0 \\ 0 & 1/b^2 \end{pmatrix}. \quad (5)$$

The inverse of the covariance matrix with respect to basis \mathcal{B} follows from a basis transformation (Eq. 3) using the rotation matrix \mathbf{R} (Eq. 4). The likelihood function (Eq. 2) then reads

$$p(\mathbf{d}|\mathbf{m}) = \frac{1}{2\pi ab} \exp \left\{ -\frac{1}{2} \left[\left(\frac{x'}{a} \right)^2 + \left(\frac{y'}{b} \right)^2 \right] \right\}, \quad (6)$$

with

$$\begin{aligned} x' &= [d_1 - d_1(\mathbf{m})] \cos \varphi + [d_2 - d_2(\mathbf{m})] \sin \varphi, \\ y' &= [d_1 - d_1(\mathbf{m})] \sin \varphi - [d_2 - d_2(\mathbf{m})] \cos \varphi, \end{aligned}$$

and (d_1, d_2) and $(d_1(\mathbf{m}), d_2(\mathbf{m}))$ are the components of \mathbf{d} and $\mathbf{d}(\mathbf{m})$, respectively.

Integrating, i.e. marginalising, the rotated Gaussian likelihood (Eq. 6) over $x \equiv d_1 - d_1(\mathbf{m})$ and $y \equiv d_2 - d_2(\mathbf{m})$ results again in Gaussian functions with standard deviations $\sigma_1 = \sqrt{a^2 \cos^2 \varphi + b^2 \sin^2 \varphi}$ and $\sigma_2 = \sqrt{b^2 \cos^2 \varphi + a^2 \sin^2 \varphi}$, respectively. These standard deviations are the conventional error bars that are typically reported in the literature, and solving the set of equations for a and b yields the desired relations for the standard deviations of the rotated Gaussian likelihood as a function of the conventional error bars of the observables and the correlation parameter,

$$\begin{aligned} a^2 &= \frac{-\sigma_1^2 \cos^2 \varphi + \sigma_2^2 \sin^2 \varphi}{\sin^4 \varphi - \cos^4 \varphi} \quad \text{and} \\ b^2 &= \frac{\sigma_1^2 \sin^2 \varphi - \sigma_2^2 \cos^2 \varphi}{\sin^4 \varphi - \cos^4 \varphi}. \end{aligned} \quad (7)$$

The above transformation from (σ_1, σ_2) to (a, b) only has unique solutions for $\varphi \neq \pm\pi/4$. In case of $\varphi = \pi/4$, $\sigma_1 = \sigma_2 \equiv \sigma$, and $\sigma^2 = (a^2 + b^2)/2$. As a consequence a and b , i.e. the covariance matrix, cannot be uniquely determined from the provided conventional uncertainties σ_1 and σ_2 , and instead the full covariance matrix has to be provided in order to use the new likelihood model. The covariance matrix is symmetric and positive-definite, i.e. its eigenvalues are real and positive: $a^2 \geq 0$ and $b^2 \geq 0$. For $a^2 \rightarrow 0$ or $b^2 \rightarrow 0$ the likelihood function describes a straight line in the d_1 - d_2 plane.

The eccentricity, e , of the shape of the likelihood function is given by ($a > b$)

$$e^2 = 1 - \frac{b^2}{a^2} = \frac{\sigma_1^2 - \sigma_2^2}{\sigma_1^2 \cos^2 \varphi - \sigma_2^2 \sin^2 \varphi}, \quad (8)$$

showing that the likelihood function has the shape of a circle for $\sigma_1 = \sigma_2$ ($\varphi \neq \pm\pi/4$) and that of an ellipse otherwise. For fixed σ_1 and σ_2 , the orientation and eccentricity of the ellipse are determined by the correlation parameter φ .

The correlation parameter describes how two observables x and y co-vary, i.e. by how much y has to change when changing x while maintaining the best possible fit/highest likelihood. Mathematically, the correlation parameter is therefore given by the derivative of y with respect to x ,

$$\tan \varphi = \frac{dy}{dx}. \quad (9)$$

For $\varphi \ll 1$ (as is the case in our examples below), Eq. (9) simplifies to $\varphi \approx dy/dx$.

Using our parametrisation, we find for the covariance of two observables p and q ,

$$\text{Cov}(p, q) = \frac{1}{2} (\sigma_p^2 - \sigma_q^2) \tan 2\varphi, \quad (10)$$

and thus for Pearson's correlation parameter

$$\rho_{pq} = \frac{1}{2} \frac{\sigma_p^2 - \sigma_q^2}{\sigma_p \sigma_q} \tan 2\varphi. \quad (11)$$

These relations show the advantage of our correlation parameter compared to other common representations of the covariance matrix, namely that φ depends only on the method with which the observables p and q have been derived, which is not the case for the covariance and Pearson's correlation parameter.

2.1.2. Three correlated observables

In Sec. 2.1.1 we derive a parametrisation of the likelihood function Eq. (2) for two correlated observables. This is already an important step because it allows us to consider correlations, for example in the HR diagram and the $T_{\text{eff}}\text{--}\log g$ plane (Kiel diagram). However, there are often three correlated observables, for example if effective temperatures, surface gravities, and luminosities of stars are known simultaneously. Analogously to Sec. 2.1.1, we now derive the relations between the conventional error bars, the eigenvalues of Σ , and the correlation parameters for three correlated observables.

The shape of the likelihood function in Eq. 2 is now an ellipsoid whose orientation in space can be described by two angles, φ and θ . Let φ be the rotation angle along the z -axis of the standard Euclidean basis \mathcal{B} and θ the rotation angle along the rotated y -axis of \mathcal{B} . Furthermore, let \mathcal{T}_1 be the basis after rotating \mathcal{B} along the z -axis by φ and \mathcal{T}_2 the basis of the reference frame after both rotations. The overall linear map describing the two rotations is then

$$\mathbf{R} \equiv \mathbf{R}(\varphi, \theta) = \mathbf{R}_y(\theta) \mathbf{R}_z(\varphi) = \mathbf{R}_z(\varphi) \mathbf{R}_y(\theta), \quad (12)$$

where $\mathbf{R}_y(\theta)$ denotes a rotation by θ along the rotated y -axis of \mathcal{B} (i.e. the y -axis of \mathcal{T}_1), $\mathbf{R}_z(\varphi)$ a rotation by φ along the z -axis of \mathcal{B} , and $\mathbf{R}_y(\theta)$ a rotation by θ along the y -axis of \mathcal{T}_1 . In the second step in Eq. (12), we have used that $\mathbf{R}_y(\theta) = \mathbf{R}_z^T(\varphi) \mathbf{R}_y(\theta) \mathbf{R}_z(\varphi)$. Overall this means that rotating first along the z -axis by φ and then by θ along the rotated y -axis of \mathcal{B} is equivalent to first rotating by θ along the y -axis and then by φ along the z -axis of \mathcal{B} . In matrix form, Eq. (12) reads

$$\mathbf{R} = \begin{pmatrix} \cos \varphi & -\sin \varphi & 0 \\ \sin \varphi & \cos \varphi & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{pmatrix}. \quad (13)$$

Using Eqs. (3) and (13), we compute the inverse of the symmetric covariance matrix (the covariance matrix with respect to \mathcal{T}_2 is diagonal and we assume the eigenvalues on the diagonal to be a^2 , b^2 , and c^2 ; cf. Eq. 5) such that the likelihood function (Eq. 2) is given by

$$p(\mathbf{d}|\mathbf{m}) = \frac{1}{(2\pi)^{3/2}abc} \exp \left\{ -\frac{1}{2} \left[\left(\frac{x'}{a} \right)^2 + \left(\frac{y'}{b} \right)^2 + \left(\frac{z'}{c} \right)^2 \right] \right\} \quad (14)$$

with

$$\begin{aligned} x' &= [d_1 - d_1(\mathbf{m})] \cos \varphi \cos \theta + [d_2 - d_2(\mathbf{m})] \cos \theta \sin \varphi \\ &\quad - [d_3 - d_3(\mathbf{m})] \sin \theta, \\ y' &= -[d_1 - d_1(\mathbf{m})] \sin \varphi + [d_2 - d_2(\mathbf{m})] \cos \varphi, \\ z' &= [d_1 - d_1(\mathbf{m})] \cos \varphi \sin \theta + [d_2 - d_2(\mathbf{m})] \sin \varphi \sin \theta \\ &\quad + [d_3 - d_3(\mathbf{m})] \cos \theta. \end{aligned}$$

In order to compute a , b , and c from the conventional uncertainties of each observable and the two rotation angles φ and θ , Eq. (14) has to be integrated. To this end, we define $x =$

$d_1 - d_1(\mathbf{m})$, $y = d_2 - d_2(\mathbf{m})$, and $z = d_3 - d_3(\mathbf{m})$ and rewrite Eq. (14) as

$$p(\mathbf{d}|\mathbf{m}) = \frac{1}{(2\pi)^{3/2}abc} \exp \left\{ -\frac{1}{2} \left[\left(\frac{\alpha_1 x + \alpha_2 y + \alpha_3 z}{a} \right)^2 + \left(\frac{\beta_1 x + \beta_2 y}{b} \right)^2 + \left(\frac{\gamma_1 x + \gamma_2 y + \gamma_3 z}{c} \right)^2 \right] \right\} \quad (15)$$

with

$$\begin{aligned} \alpha_1 &= \cos \varphi \cos \theta, \\ \alpha_2 &= \cos \theta \sin \varphi, \\ \alpha_3 &= -\sin \theta, \\ \beta_1 &= -\sin \varphi, \\ \beta_2 &= \cos \varphi, \\ \gamma_1 &= \cos \varphi \sin \theta, \\ \gamma_2 &= \sin \varphi \sin \theta, \\ \gamma_3 &= \cos \theta. \end{aligned} \quad (16)$$

With these definitions we find

$$\begin{aligned} L(x) &\equiv \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} p(\mathbf{d}|\mathbf{m}) dy dz = \frac{1}{\sqrt{2\pi}\sigma_1} \exp \left[-\frac{1}{2} \left(\frac{x}{\sigma_1} \right)^2 \right], \\ L(y) &\equiv \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} p(\mathbf{d}|\mathbf{m}) dx dz = \frac{1}{\sqrt{2\pi}\sigma_2} \exp \left[-\frac{1}{2} \left(\frac{y}{\sigma_2} \right)^2 \right], \\ L(z) &\equiv \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} p(\mathbf{d}|\mathbf{m}) dx dy = \frac{1}{\sqrt{2\pi}\sigma_3} \exp \left[-\frac{1}{2} \left(\frac{z}{\sigma_3} \right)^2 \right], \end{aligned}$$

which are Gaussian functions with standard deviations

$$\begin{aligned} \sigma_1^2 &= a^2 \beta_2^2 \gamma_3^2 + b^2 (\alpha_3 \gamma_2 - \alpha_2 \gamma_3)^2 + c^2 \alpha_3^2 \beta_2^2, \\ \sigma_2^2 &= a^2 \beta_1^2 \gamma_3^2 + b^2 (\alpha_3 \gamma_1 - \alpha_1 \gamma_3)^2 + c^2 \alpha_3^2 \beta_1^2, \\ \sigma_3^2 &= a^2 (\beta_1 \gamma_2 - \beta_2 \gamma_1)^2 + b^2 (\alpha_1 \gamma_2 - \alpha_2 \gamma_1)^2 + c^2 (\alpha_1 \beta_2 - \alpha_2 \beta_1)^2. \end{aligned}$$

Using the definitions from Eqs. (16) and solving for the semi-major and -minor axes a , b , and c of the ellipsoid, we arrive at the desired relation between the conventional uncertainties of each observable and the semi-major and -minor axes,

$$\begin{aligned} a^2 &= \frac{\cos^2 \theta (\sigma_1^2 \cos^2 \varphi - \sigma_2^2 \sin^2 \varphi) - \sigma_3^2 \sin^2 \theta (\cos^2 \varphi - \sin^2 \varphi)}{(\cos^2 \varphi - \sin^2 \varphi) (\cos^2 \theta - \sin^2 \theta)}, \\ b^2 &= \frac{\sigma_1^2 \sin^2 \varphi - \sigma_2^2 \cos^2 \varphi}{\sin^4 \varphi - \cos^4 \varphi}, \\ c^2 &= \frac{\sigma_3^2 \cos^2 \theta (\cos^2 \varphi - \sin^2 \varphi) - \sin^2 \theta (\sigma_1^2 \cos^2 \varphi - \sigma_2^2 \sin^2 \varphi)}{(\cos^2 \varphi - \sin^2 \varphi) (\cos^2 \theta - \sin^2 \theta)}. \end{aligned}$$

As in the case of two correlated observables (Sec. 2.1.1), the full specification of our new likelihood model from conventional uncertainties is only possible if $\varphi \neq \pm\pi/4$ and $\theta \neq \pm\pi/4$. Furthermore, the combinations of σ_1 , σ_2 , σ_3 , φ , and θ have to be such that $a^2 > 0$, $b^2 > 0$, and $c^2 > 0$ because Σ is symmetric and positive-definite.

We define the correlation parameters to describe correlations of two observables with respect to \mathcal{B} . With the definitions of φ and θ as rotations along the z -axis and the rotated y -axis, φ already describes the correlation between two observables in the x - y plane, but θ does not describe the correlation in either the x - z or y - z plane of \mathcal{B} , but in the rotated x - z plane of \mathcal{B} . Hence, if ξ is the correlation parameter of two observables in the x - z plane, θ and ξ are related through

$$\tan \theta = \tan \xi \cos \varphi. \quad (17)$$

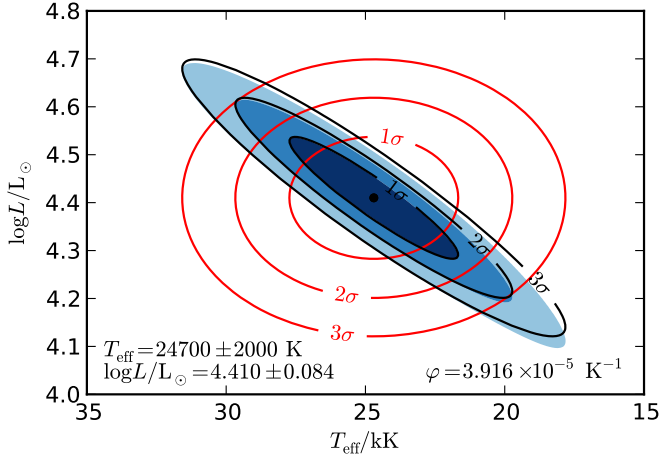


Fig. 2. Comparison of the reference likelihood function with two approximations that neglect and account for correlations. The 1σ , 2σ , and 3σ areas of the reference likelihood function in the HR diagram are shown by the shaded areas. The red contours show the same confidence levels of a likelihood model that neglects the correlation in luminosity and effective temperature, and the black contours show our new likelihood model for a correlation parameter of $\varphi = 3.916 \times 10^{-5} \text{ K}^{-1}$. The marginalised observables are $T_{\text{eff}} = 24700 \pm 2000 \text{ K}$ and $\log L/L_{\odot} = 4.410 \pm 0.084$ for all three likelihood models.

2.2. Determination of correlation parameters

Whether observables co-vary or not depends solely on how they are determined. Observables can also be fully independent of each other, which is the case if it is possible to determine one parameter without needing to know the others or if observables are inferred from independent methods. On the contrary, it may be that the determination of parameters depends (strongly) on knowing the others, in which case the correlation between the observables may add valuable information when comparing them to stellar models.

Determining correlation parameters can be complex. In all probability, the easiest situation is if observables are related to each other via an analytical function that is used to derive some observables from others. In this case the correlation parameters follow from Eq. (9). We discuss such an example in Sec. 2.3.

If there is no analytical relation, the correlation parameter can be found by computing the covariance of those observables that will be matched against stellar models. To this end, the relevant parameter space has to be sampled such that a (multi-dimensional) probability distribution for the relevant observables can be computed. The correlation parameter can then be determined by fitting our new likelihood model for two correlated observables (Eq. 6) to the corresponding marginalised probability distribution. This method can always be applied, but may be computationally expensive. In Sec. 2.4 we discuss such an example and also show how it is possible to determine the correlation parameter more easily.

2.3. Correlations in the Hertzsprung–Russell diagram

Our new likelihood model can be applied to any kind of star in any situation as long as the correlations are known. In the following, we turn to massive, main-sequence stars because such models are already available in BONNSAI. We consider the situa-

tion that the effective temperature of a star is known, for example from fitting its spectral energy distribution or using a spectral type calibration, the apparent V-band magnitude from photometric observations, and the distance from parallax measurements. The luminosity of the star can then be computed using the bolometric correction calibration of Flower (1996). We choose the bolometric correction of Flower (1996) for demonstration purposes because it is valid over a wide range of stellar temperatures, from M- to O-type stars. For this example, we assume that the effective temperature is $T_{\text{eff}} = 24700 \pm 2000 \text{ K}$, the distance to the star $d = 151 \pm 5 \text{ pc}$, the apparent V-band magnitude $m_V = 1.97 \pm 0.01$, and that there is only negligible extinction (these observables, except for the large uncertainty in effective temperature, are characteristic of the magnetic B-type star HD 44743; Fossati et al. 2015). The luminosity of the star is then given by

$$\log L/L_{\odot} = 0.4 \left(M_{V,\odot} - m_V + 5 \log \frac{d}{\text{pc}} - 5 - BC_V(T_{\text{eff}}) \right), \quad (18)$$

where $BC_V(T_{\text{eff}})$ is the bolometric correction and $M_{V,\odot}$ the absolute V-band magnitude of the Sun. This equation shows that the logarithmic luminosity co-varies with apparent magnitude, distance, and effective temperature. On a more fundamental level, the reason why luminosity and effective temperature co-vary is that both quantities are given by the total flux of stars such that the Stefan–Boltzmann law holds, $L = 4\pi R^2 \sigma T_{\text{eff}}^4$ with R being the stellar radius and σ the Stefan–Boltzmann constant. We now consider the case where the luminosity and effective temperature are matched against stellar models to infer fundamental stellar parameters for this star. We are therefore only interested in the covariance of luminosity and effective temperature, and show in Fig. 2 the likelihood function computed from Eq. 18 using the above stellar parameters and assuming Gaussian uncertainties; the marginalised luminosity is then $\log L/L_{\odot} = 4.410 \pm 0.084$. This likelihood model makes the fewest number of assumptions on the data and is therefore the most accurate model discussed here. In order to easily distinguish between the different likelihood models, we call it the reference likelihood model from here on. We further show the old Gaussian likelihood model that neglects correlations for $T_{\text{eff}} = 24700 \pm 2000 \text{ K}$ and $\log L/L_{\odot} = 4.410 \pm 0.084$. The two likelihood functions differ significantly.

Our new likelihood model is parametrised by the same luminosity and effective temperature as the old model with the addition of a correlation parameter φ that describes how the luminosity co-varies with effective temperature. Here, $\varphi \ll 1$, and applying Eq. (9) we find

$$\varphi(T_{\text{eff}}) = \frac{d \log L/L_{\odot}}{dT_{\text{eff}}} = -0.4 \frac{dBC}{dT_{\text{eff}}}. \quad (19)$$

Using the analytic fit of the bolometric correction as a function of effective temperature from Flower (1996), i.e. taking the derivative of the fit of the bolometric correction with respect to effective temperature and evaluating the derivative at $T_{\text{eff}} = 24700 \text{ K}$, we find $\varphi = 3.916 \times 10^{-5} \text{ K}^{-1}$. It is evident from Fig. 3 that the new likelihood model represents the reference likelihood well and much better than the old model.

Because of the different shape and orientation between the new likelihood model and the old, certain combinations of effective temperature and luminosity can be excluded with more than 3σ confidence, which would be considered possible within 1σ when neglecting the correlations. Given that different combinations of effective temperature and luminosity correspond to

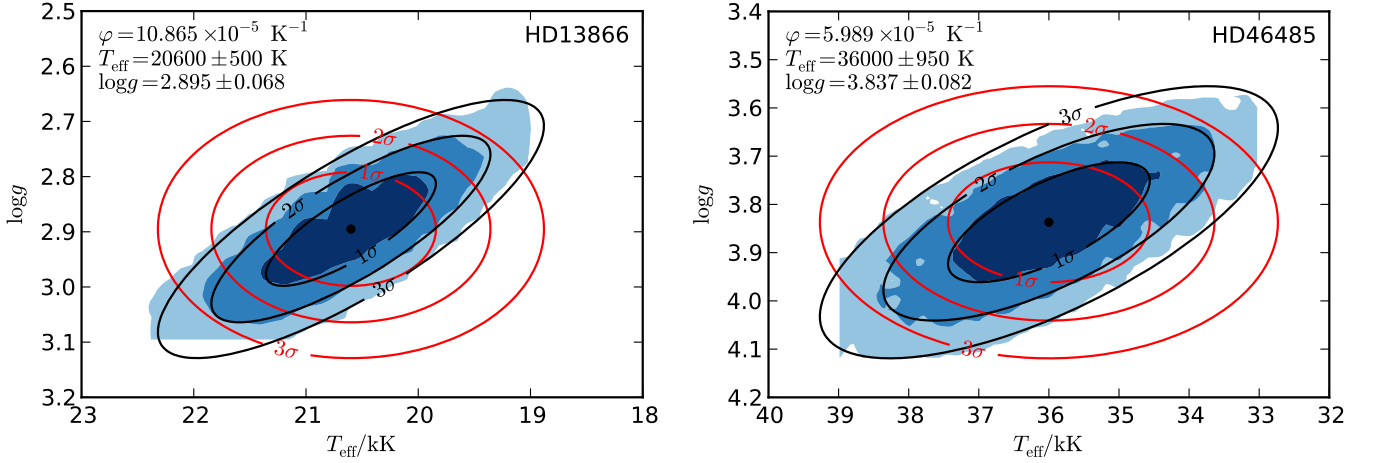


Fig. 3. As in Fig. 2 but for (a) HD 13866 and (b) HD 46485 in the Kiel diagram. The probability maps have been derived by fitting FASTWIND atmosphere models to observed spectra.

different combinations of the fundamental stellar parameters initial mass, age, and initial rotational velocity, this will influence the inference of these parameters (cf. Fig. 4 and Sec. 3 where we show that the most likely stellar parameters and the inferred uncertainties can change significantly).

Although the new likelihood model approximates the reference likelihood far better than a likelihood model that neglects correlations, the new approximation is not perfect. The correlation parameter varies with effective temperature leading to small deviations from the reference likelihood visible only in the $2\text{--}3\sigma$ regions. Halving the large uncertainty in effective temperature makes the deviations disappear nearly completely. We therefore conclude that the new likelihood model is an adequate approximation in this case.

2.4. Correlations in the Kiel diagram

The second example of correlations concerns effective temperatures and surface gravities that are deduced from fitting synthetic spectra of atmosphere models to observed ones. The correlation between surface gravity and effective temperature from fitting the spectra of stars depends on the details of the fitting process and the applied methods (see below). For example, in OB stars the correlation is such that hotter temperatures require larger gravities to fit spectra similarly well, whereas the opposite is true in cooler stars. In Wolf–Rayet stars, the photosphere is formed in the wind and the spectral lines are therefore not sensitive to the surface gravity. Hence, effective temperature and surface gravity are uncorrelated in these stars (in fact, the surface gravity remains mostly unconstrained). These examples show that the correlations of observables depend only on how observables are determined (cf. Sec. 2.2).

In Fig. 3 we show how effective temperature and surface gravity co-vary when modelling the spectra of HD13866 and HD46485, observed within the IACOB project (Simón-Díaz et al. 2011, 2015), with the stellar atmospheric code FASTWIND (Santolaya-Rey et al. 1997; Puls et al. 2005). The probability distributions in Fig. 3 have been computed in two steps. First, the main optical transitions in the observed spectra were compared to a pre-computed grid of FASTWIND stellar atmosphere models presented in Castro et al. (2012) to find the best-fitting spectroscopic stellar parameters (for further details on the technique and the lines used, see Castro et al. 2012). Second, sub-grids of

atmosphere models have been computed around the best-fitting effective temperature and surface gravity to properly resolve the probability distributions such that the correlation parameter can be reliably determined. The finer sub-grids have a resolution of $\Delta T_{\text{eff}} = 250 \text{ K}$ and $\Delta \log g = 0.025 \text{ dex}$ and the probability map is a bit patchy because of this finite grid in atmosphere models. From the probability maps we find correlation parameters of $10.865 \times 10^{-5} \text{ K}^{-1}$ and $5.989 \times 10^{-5} \text{ K}^{-1}$ for HD 13866 and HD 46485, respectively. In this example, the correlation is stronger in the cooler OB star.

As in Sec. 2.3 for a star in the HR diagram, it is evident that correlations can significantly change the shape and orientation of the likelihood function. The new likelihood model nicely matches the probability maps and properly reproduces the different confidence regions. A comparison with the old likelihood model reveals the need for taking correlations properly into account to avoid biases when determining fundamental stellar parameters such as mass and age.

Determining the correlation parameter is not trivial. As discussed in Sec. 2.2, a viable method to obtain the correlation parameter is to fit our likelihood model to detailed probability maps such as those in Fig. 3; this procedure can be computationally expensive. An alternative, computationally less expensive method works as follows: First, the best-fitting spectroscopic parameters are determined by varying those parameters that significantly influence the diagnostic lines. Second, the best-fitting effective temperature is changed by δT_{eff} and $\log g$ is varied by $\delta \log g$ until the deviation of the synthetic spectrum from the observed spectrum is minimised again. During this step, all other parameters can be kept constant at the best-fitting values and the correlation parameter follows from $\varphi = \delta \log g / \delta T_{\text{eff}}$ because $\varphi \ll 1$ (Eq. 9). It is advisable to repeat the steps for several values of δT_{eff} to obtain a more robust correlation parameter. Both procedures can of course be applied to any stellar parameter and not only effective temperature and surface gravity.

The absolute value of the correlation parameter of effective temperature and surface gravity depends on a variety of factors; it depends on which atmospheric lines are used in the analysis (e.g. hydrogen–helium lines or hydrogen–helium–silicon lines) and which weight is given to the diagnostic lines. For example, the surface gravity is strongly constrained by the wings of the Balmer lines and some fitting methods require always fitting the Balmer wings well while allowing for more freedom in other

lines. The Balmer wings and thus the derived surface gravity further depend on the signal-to-noise ratio (S/N), i.e. the correlation parameter will also be a function of S/N. Finally, the correlation parameter depends on the atmospheric parameters varied in the fitting process because all parameters have the potential to change the shape and strengths of the main-diagnostic lines. For example, allowing for variations in the helium abundances obviously influences the helium lines and thus the correlation parameter; the wind-Q parameter also influences some helium lines (e.g. Puls et al. 1996; Kudritzki & Puls 2000).

2.5. Normalised correlation parameter

The deviation of the new likelihood model from a likelihood model that neglects correlations depends on the absolute value of the correlation parameter and it also depends on the conventional error bars. If a quantity y co-varies strongly with x but the conventional uncertainty of x is much smaller than that of y , the new likelihood model deviates only slightly from a likelihood that neglects correlations. Similarly, the deviation of the likelihoods can be significant if y co-varies only weakly with x but the uncertainty of x is much larger than that of y . The importance of the correlations therefore depends on the interplay between the correlation parameter and the error bars (cf. Eq. 8).

In order to judge the importance of correlations, it is useful to define a normalised correlation parameter, which is the correlation parameter introduced in Sec. 2.1.1 divided by the ratio of the corresponding 1σ uncertainties:

$$\hat{\varphi} = \varphi / \frac{\sigma_2}{\sigma_1} = \frac{dy}{dx} / \frac{\sigma_2}{\sigma_1}. \quad (20)$$

This parameter is a measure of the importance of correlations with $\hat{\varphi} \rightarrow 0$ indicating no influence and $\hat{\varphi} \rightarrow 1$ maximum influence on the shape of the likelihood (in the former case, the likelihood function equals a likelihood that neglects correlations, whereas it converges to a straight line in the latter).

A geometric interpretation of the normalised correlation parameter may be obtained by considering $\varphi \ll 1$. This applies for correlations between effective temperature and logarithmic luminosity and between effective temperature and logarithmic surface gravity discussed in Secs. 2.3 and 2.4, respectively. The ratio of the 1σ areas of the likelihood function accounting for and neglecting correlations is then

$$\frac{\pi ab}{\pi \sigma_1 \sigma_2} \approx \sqrt{1 - \left(1 + \sigma_2^4 / \sigma_1^4\right) \hat{\varphi}^2} \approx \sqrt{1 - \hat{\varphi}^2}. \quad (21)$$

The last approximation holds if $\sigma_2^4 / \sigma_1^4 \ll 1$, which is the case for the mentioned examples. In these cases, the normalised correlation parameter relates directly to the relative change in the 1σ area of the likelihood function. This interpretation also allows us to quantify the importance of correlations for the likelihood function: a larger than 5% (10%) reduction of the 1σ area corresponds to $\hat{\varphi} \gtrsim 0.31$ ($\hat{\varphi} \gtrsim 0.43$).

For the example of star HD 44743 in the HR diagram (Sec. 2.3), the normalised correlation parameter is $\hat{\varphi} = 0.93$, indicating a strong correlation that has large consequences for the likelihood model. Indeed, the area of the 1σ region is, in accordance with Eq. (21), reduced by about 63% when accounting for correlations (Fig. 2). In the Kiel diagram, the normalised correlation parameters are $\hat{\varphi} = 0.8$ and $\hat{\varphi} = 0.7$ for HD 13866 and HD 46485, respectively (Sec. 2.4), leading to a reduction of about 29% and 40% of the 1σ area when accounting for correlations (Fig. 3).

The absolute correlation parameters of HD 13866 and HD 46485 in the Kiel diagram are larger than that of HD 44743 in the HR diagram (Secs. 2.3 and 2.4). However, the shape of the likelihood of HD 44743 is more elongated (Fig. 2), i.e. the correlation is more important, than for HD 13866 and HD 46485 (Fig. 3). This shows the interplay between the correlation parameter and the 1σ uncertainties, and highlights the value of the normalised correlation parameter in providing a degree for the strength of correlations.

3. Results

As shown in Secs. 2.3 and 2.4, likelihood models that neglect and take correlations into account can cover significantly different regions of the parameter space and will therefore affect the determination of fundamental stellar parameters such as mass and age. To this end we investigate changes in the inferred initial masses and ages for a sample of nearly 500 mock stars distributed all over the HR and Kiel diagrams. The mock stars are taken from the Milky Way stellar models of Brott et al. (2011) and initial masses and ages are distributed in such a way to achieve a good coverage of the HR and Kiel diagrams. We choose 1σ uncertainties of 1000 K, 0.1 dex, and 0.1 dex in effective temperature, logarithmic luminosity, and logarithmic surface gravity, respectively. These uncertainties are characteristic values for spectroscopically derived stellar parameters. The correlation parameter of effective temperature and luminosity in Sec. 2.3 is about $0.4 \times 10^{-4} \text{ K}^{-1}$, resulting in a normalised correlation parameter of about 0.4 for our choice of the 1σ uncertainties. We choose a normalised correlation parameter of 0.8 for effective temperature and surface gravity reminiscent of those found for HD 13866 and HD 46485 in Sec. 2.4. We thus introduce one case of modest and one case of notable correlation (reductions in the 1σ area of about 8% and 40%, respectively), and therefore expect that the most likely values and precisions are more strongly affected for stars with known effective temperatures and gravities than for stars with known effective temperatures and luminosities. A different choice of error bars (e.g. smaller uncertainties on luminosities) may lead to the opposite situation. We therefore caution that the quantitative results presented here are only valid for our particular choice of correlations and error bars.

In our analysis, we use a Salpeter initial mass function (Salpeter 1955) as initial mass prior distribution, a uniform age prior distribution, and a Gaussian initial rotational velocity prior distribution with a mean of 100 km s^{-1} and full width at half maximum of 250 km s^{-1} (the initial rotational velocity prior distribution is thus reminiscent of the observed distribution of rotational velocities of B-type stars in the Milky Way, Hunter et al. 2008). BONNSAI samples all those Milky Way stellar models from Brott et al. (2011) from a pre-computed database that are within 5σ of the observations. Given that the mock stars are taken from the same stellar models and the chosen size of the observational error bars, the posterior-predictive check (goodness-of-fit test) and resolution test are passed in each case.

We first discuss the qualitative changes in inferred masses and ages because of correlations (Sec. 3.1) before presenting quantitative results for the precisions (Sec. 3.2) and most likely (Sec. 3.3) initial masses and ages. Also, the choice of prior distributions influences the inference of mass and age, and we examine this in Appendix A.

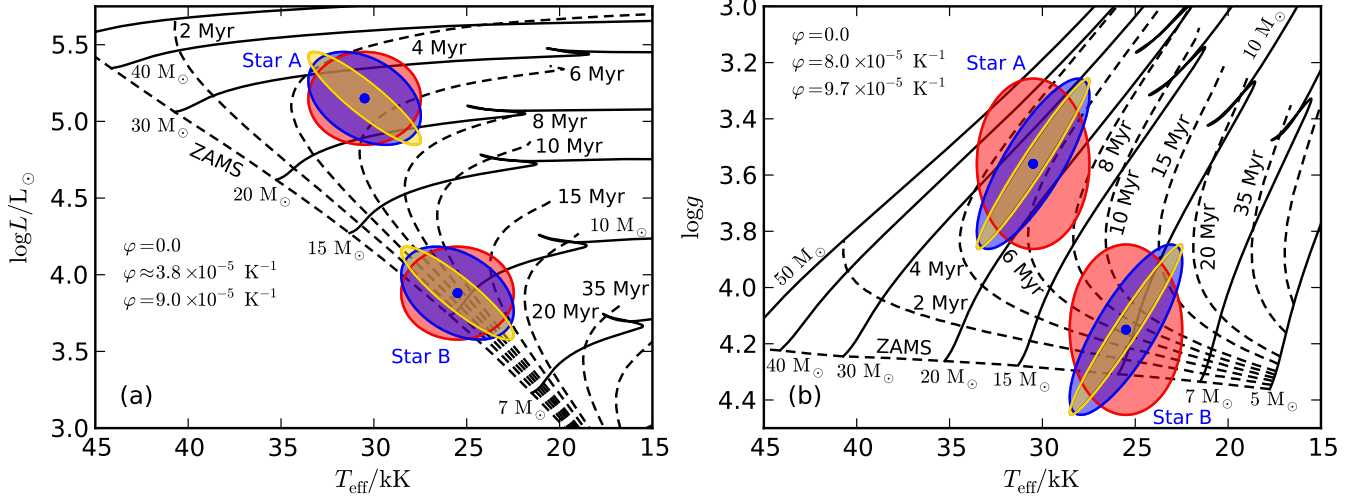


Fig. 4. Mock stars, Star A and B, in the HR diagram, panel (a), and Kiel diagram, panel (b). The error ellipses indicate 1σ confidence regions and we increased our standard error bars by a factor of 2 for illustration purposes. The red ellipses are for the case of uncorrelated observables and the blue ellipses for correlations as described in Secs. 2.3 and 2.4, and applied in Secs. 3.2 and 3.3. For illustration purposes, the golden ellipses show very strong correlations. The solid lines are the non-rotating, Milky Way metallicity stellar tracks of Brott et al. (2011) for masses ranging from 5 to $50 M_{\odot}$, and the dashed lines are the corresponding isochrones from 0 to 50 Myr.

3.1. Qualitative discussion of the influence of correlations on inferred stellar parameters

In order to illustrate the expected changes in inferred stellar parameters, we show the position of two of the approximately 500 mock stars, Star A and B, in the HR and Kiel diagram (Fig. 4). In Fig. 4 we show the 1σ contours for (i) no correlations, (ii) the correlations used in the quantitative analysis in Secs. 3.2 and 3.3 (see above), and (iii) very strong correlations. The error ellipses are not rotated for non-correlations and turn into straight lines for maximum correlations.

Because of the correlations, the error ellipses change their orientation relative to the stellar tracks and isochrones. If an error ellipse is parallel to the stellar tracks (isochrones), the 1σ region covers fewer different masses (ages) than the error ellipse in the case of uncorrelated observables, meaning that the precision with which the mass (age) of a star is determined increases. The precision can also decrease if the error ellipse happens to be perpendicular to the stellar tracks (isochrones) such that a wider range of masses (ages) is covered. For Star A, the stellar tracks and isochrones are highly inclined with respect to the semi-major axis of the error ellipse in the HR diagram and about parallel in the Kiel diagram. Consequently, mass and age can be determined to a lower (higher) precision from the HR (Kiel) diagram when taking correlations into account. For star B in the HR diagram the age can be determined with a higher precision, while the mass is less precisely known when taking correlations into account. The situation is opposite in the Kiel diagram: the mass can now be determined to a higher precision, while age only to a lower precision.

Not only does the precision of inferred stellar parameters change, but also their most likely values. The most likely model takes the model density of stars because of different stellar lifetimes in various regions of the parameter space and prior knowledge such as the initial mass function into account. The most likely values of inferred stellar parameters are unaffected by correlations only if the model density is homogeneous and uniform prior distributions are applied. However, the model den-

sity is not homogeneous, but increases towards less massive and younger stars, and initial-mass prior distributions usually favour lower mass stars. In the HR diagram the model density is therefore highest towards hotter effective temperatures and less luminous stars, and in the Kiel diagram towards cooler effective temperatures and higher gravities. The correlations discussed in Secs. 2.3 and 2.4 are exactly the opposite: hotter temperatures require brighter luminosities in the HR diagram and larger gravities in the Kiel diagram. Hence, the correlations work against the tendency that the most likely initial mass and age is found towards the highest model densities. Neglecting correlations can therefore bias inferred stellar parameters.

We want to stress that no matter whether it is the precisions or the inferred most likely values that change when taking correlations into account, the reliability and robustness of the inferred fundamental stellar parameters and their uncertainties always improve. This may be counterintuitive, especially when parameters can only be determined to a lower precision when accounting for correlations, but it simply means that inferred uncertainties are underestimated otherwise.

3.2. Precision of inferred initial masses and ages

3.2.1. Stars in the HR diagram

We explore the precision, p , with which the initial mass and age of our mock stars can be determined when neglecting correlations and then compare it to the case when taking correlations into account. The precision is defined as the ratio of the 1σ uncertainty and the mode (most likely) value of the posterior probability distributions; it is shown in the left panels (a) and (c) of Fig. 5. In the right panels (b) and (d) of the same figure, we show the ratio of precisions taking correlations into account and neglecting them, $r = p_{\text{corr}}/p_{\text{no corr}}$. A ratio of $r > 1$ indicates a lowering and a ratio of $r < 1$ an improvement in the precisions when taking correlations into account. The ratio is defined such that the product of the ratios in panels (b) and (d), and the pre-

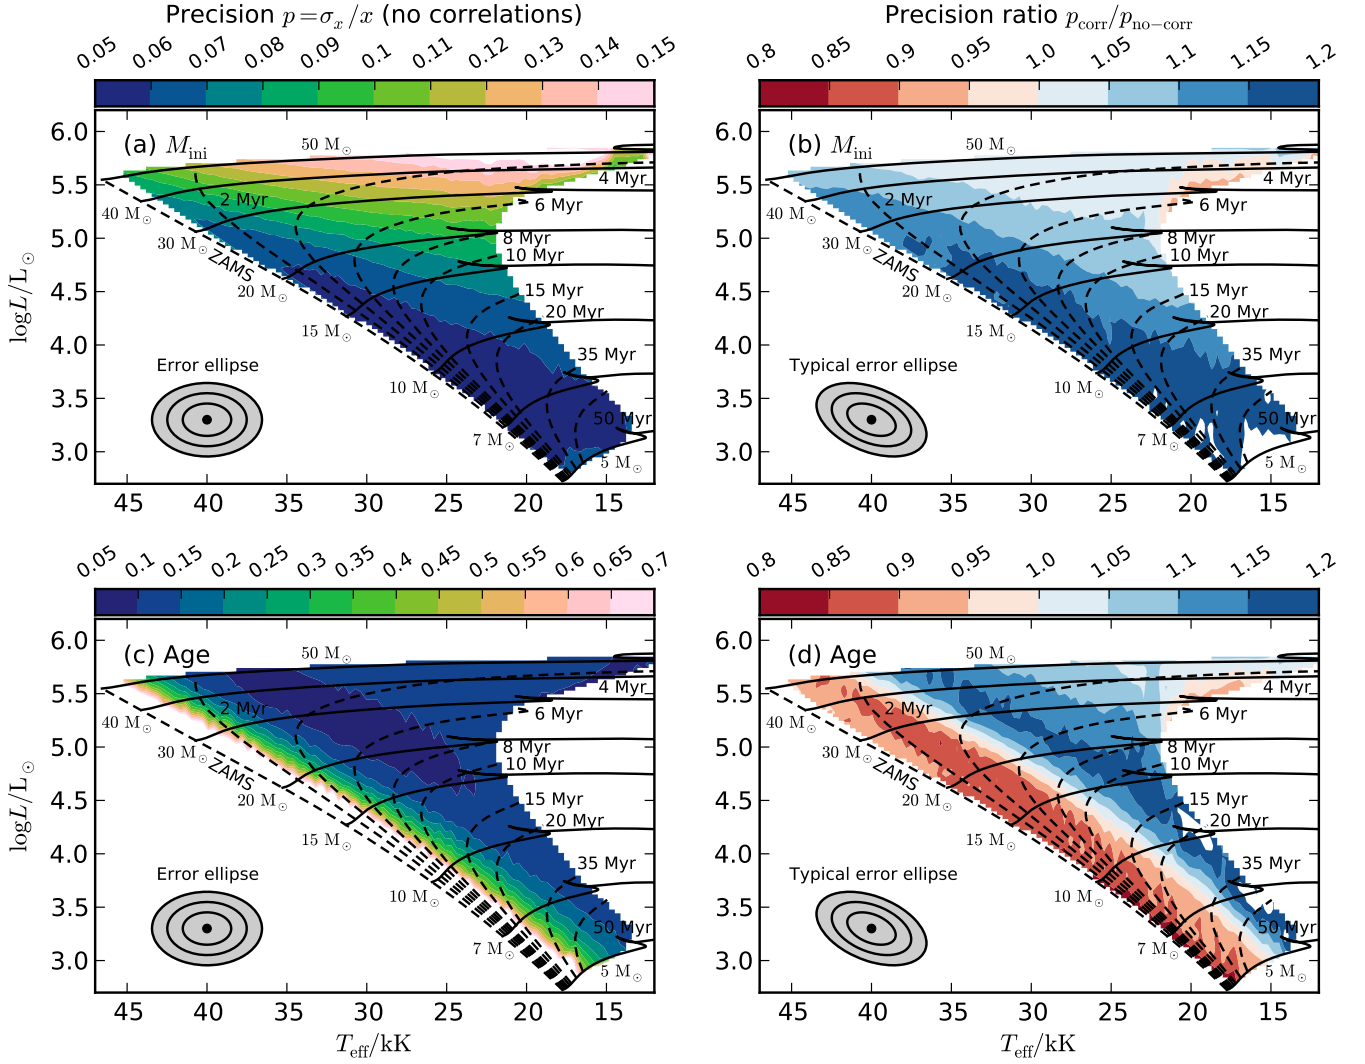


Fig. 5. Precisions $p = \sigma_x/x$ (left panels) and precision ratios $p_{\text{corr}}/p_{\text{no-corr}}$ (right panels) of determined initial masses (upper panels) and ages (lower panels) of our mock stars in the HR diagram (here, σ_x and x are the 1σ error bars and most likely values of initial mass and age, respectively, and p_{corr} and $p_{\text{no-corr}}$ denote the precisions when considering and neglecting correlations, respectively). The colour-coding in the left panels (a) and (c) show the precision neglecting correlations and the colours in the right panels (b) and (d) the ratio of the precisions taking correlations into account. The 1σ uncertainties of the mock stars are 1000 K in effective temperature and 0.1 dex in luminosity, and the resulting 1σ , 2σ , and 3σ contours of the likelihood function are shown in the lower left corners (a correlation parameter of $4 \times 10^{-5} \text{ K}^{-1}$ is used for illustration purposes in panels b and d). The plotted stellar tracks and isochrones are Milky Way metallicity, non-rotating models from [Brott et al. \(2011\)](#).

cisions in panels (a) and (c) directly give the precisions when taking correlations into account.

Neglecting correlations, the initial mass and age can be determined up to a precision of about 5% for the given uncertainties in effective temperature and luminosity. The precision scales roughly linearly with the error bars, i.e. halving the uncertainties of effective temperature and luminosity means that masses and ages can be determined two times more precisely. Masses and ages can be determined to the highest precision wherever the spacing between stellar tracks and isochrones is largest, giving rise to the gradients in Figs. 5a and 5c. Formally, the age precision diverges for stars approaching the zero age main sequence (ZAMS). We show precisions down to 70%, resulting in uncoloured areas close to the ZAMS.

As illustrated in Fig. 4, taking correlations into account changes the precision with which mass and age can be deter-

mined depending on the relative orientation of the likelihood function to the stellar tracks and isochrones. The precision in mass always decreases because the error ellipses are highly inclined with respect to the stellar tracks such that the rotated error ellipses cover a wider range of masses; the decrease can be up to 20% (Figs. 5b and 5d). The precision in age increases for stars close to the ZAMS because the orientation of the error ellipse is almost parallel to the isochrones and decreases once the isochrones bend over and become more parallel to the stellar tracks. The bending of the isochrones and the change in the precisions are nicely illustrated in Fig. 5d and amount to up to about $\pm 20\%$.

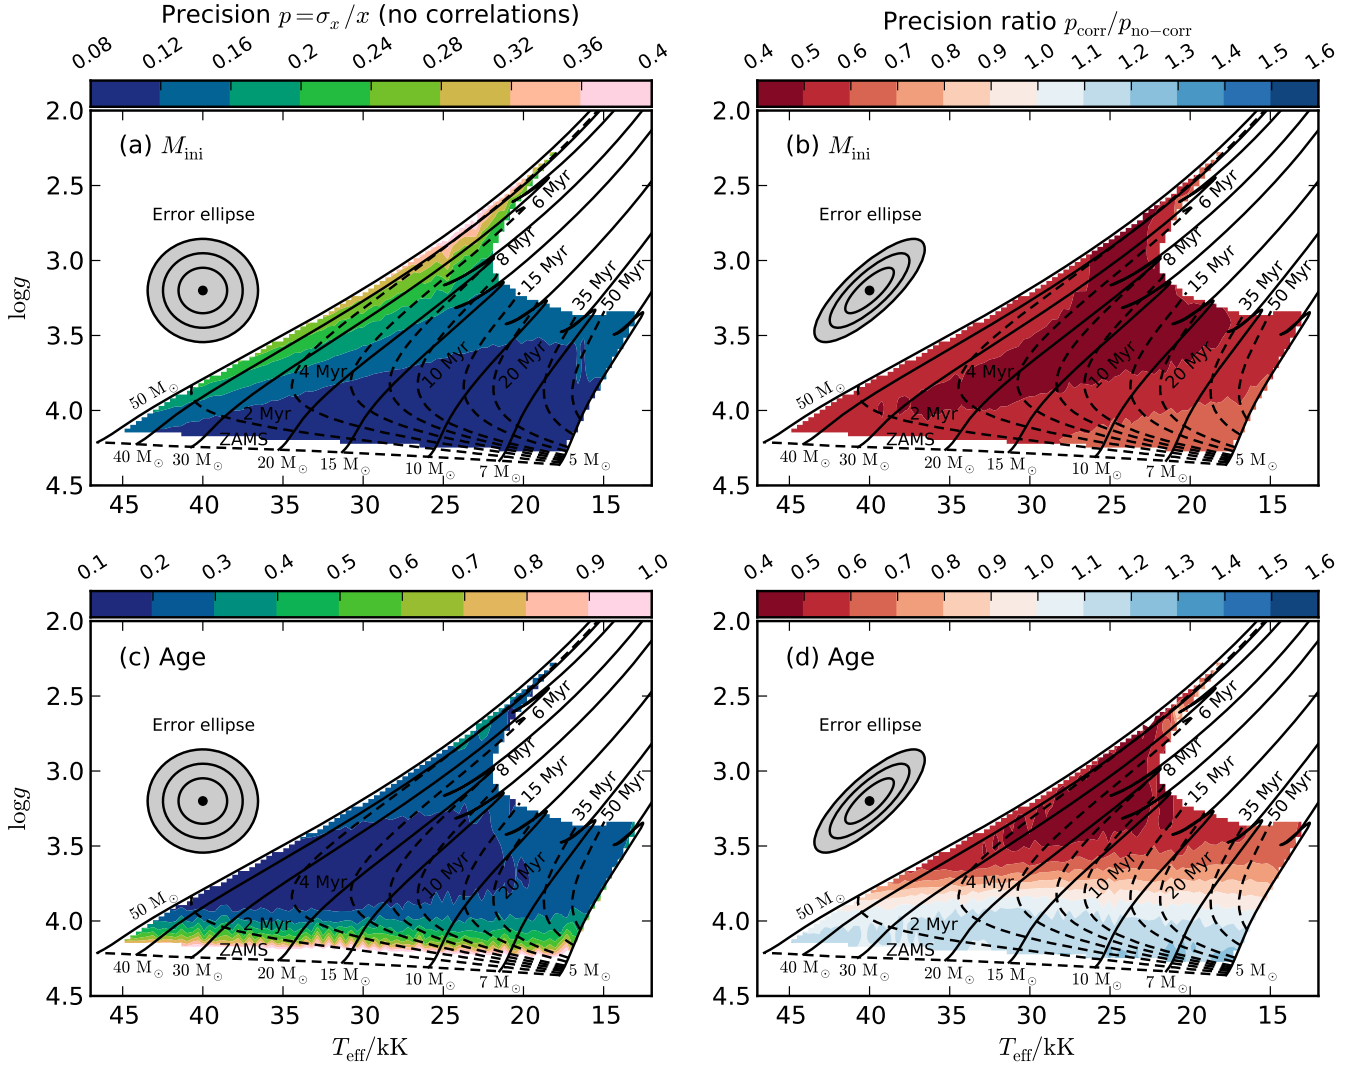


Fig. 6. Same as Fig. 5 but for stars in the Kiel diagram. The 1σ uncertainties are 1000 K in effective temperature and 0.1 dex in logarithmic surface gravity, and a correlation parameter of $0.8 \times 10^{-4} \text{ K}^{-1}$ is assumed for all stars (cf. Sec. 2.4). The error ellipses show 1σ , 2σ , and 3σ uncertainty contours.

3.2.2. Stars in the Kiel diagram

Analogously to Sec. 3.2.1, we study the precision with which masses and ages can be determined from the position of stars in the Kiel diagram. Neglecting correlations, masses, and ages can be determined with a precision of up to about 10% (Fig. 6). As discussed in Sec. 3.2.1, the precision in mass and age is the best for stars in those areas of the Kiel diagram where the spacing between stellar tracks and isochrones is the largest.

Unlike stars in the HR diagram, the precision with which masses can be determined always improves when considering correlations. For our choice of the error bars and the correlation parameter, the precision in mass improves by up to 60%, i.e. masses can be determined almost twice as precisely. This also means that the luminosities of stars can be predicted more precisely and hence also the spectroscopic distance. Furthermore, the age can be determined significantly more precisely (up to 60%) in those areas in the Kiel diagram where the isochrones tend to be parallel to the error ellipse. Close to the ZAMS the isochrones are almost perpendicular to the rotated error ellipses and the precision in age lowers. Overall, the changes in precision

in mass and age are larger for our stars in the Kiel diagram than for stars in the HR diagram because of our choice of stronger correlations in effective temperature and surface gravity rather than of effective temperature and luminosity.

3.2.3. Combined HR and Kiel diagrams

We now match effective temperatures, surface gravities, and luminosities simultaneously to stellar models while applying the same correlations as before. The precisions with which mass and age can be determined are shown in Fig. 7. In general it is expected that the mass and age precisions improve (or at least remain the same) when matching all three observables against stellar models and not just two observables because there is more information to better constrain fundamental stellar parameters. However, the improvements may not be large and are only significant whenever the third observable adds valuable information. In Fig. 4 the chosen error ellipse for Star A and B in the Kiel diagram span a wider range of initial masses and ages than in the HR diagram. This is reflected in the mass and age precisions from stars in the HR and Kiel diagrams where we find that both

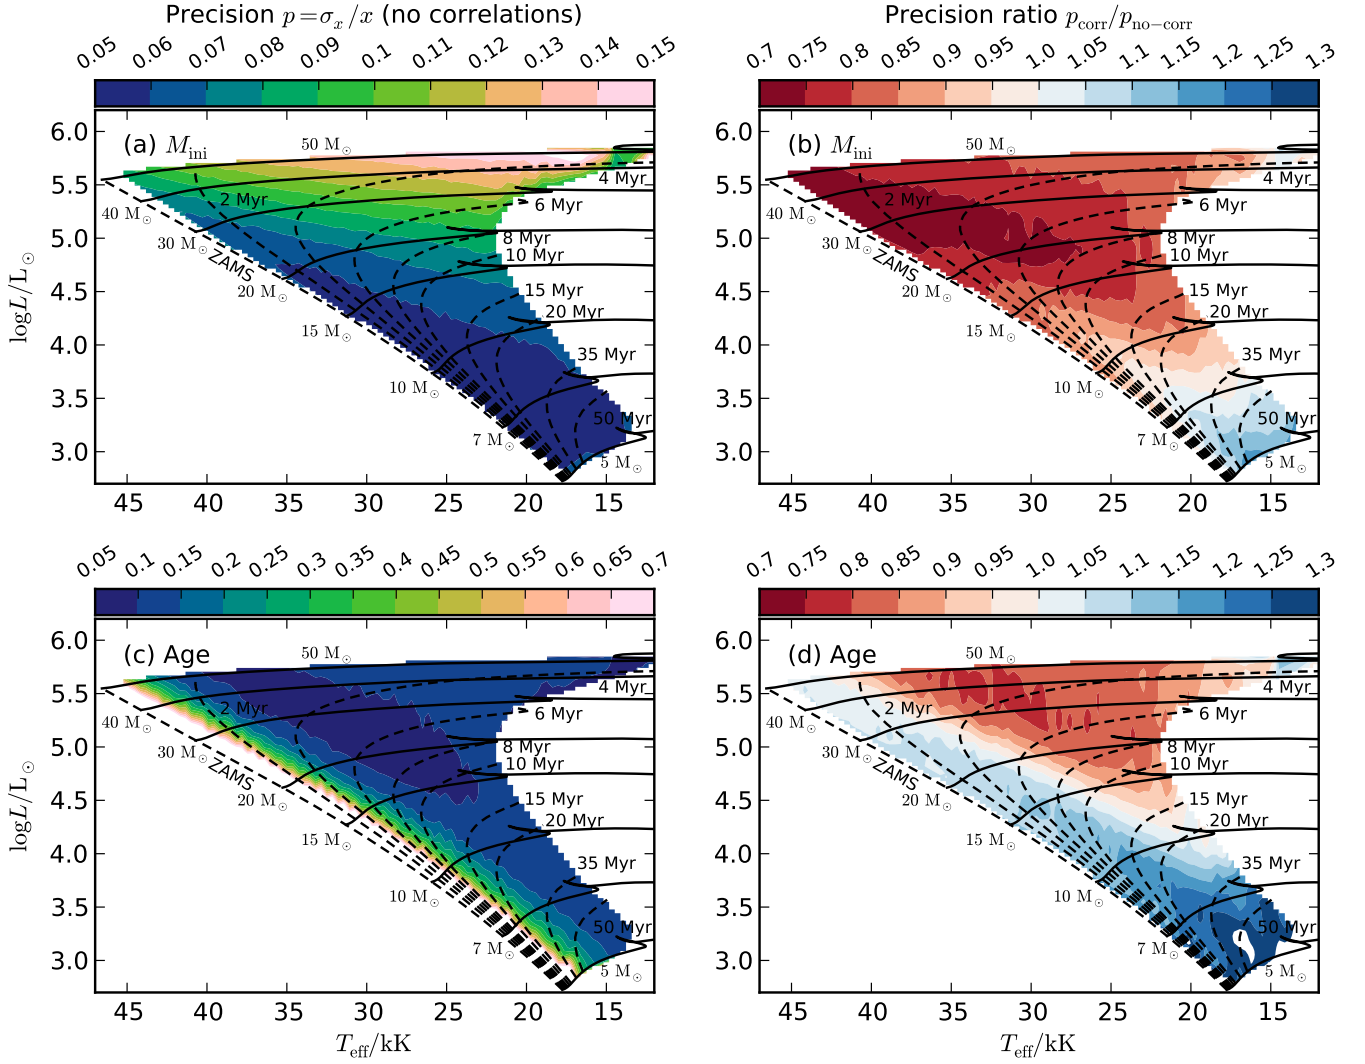


Fig. 7. Same as Figs. 5 and 6 but for stars with known positions in the HR and Kiel diagrams. The 1σ uncertainties are 1000 K in effective temperature, 0.1 dex in logarithmic luminosity, and 0.1 dex in logarithmic surface gravity. The correlation parameters are the same as in Secs. 3.2.1 and 3.2.2.

mass and age can be inferred to a higher precision from the HR diagram because of our choice of error bars (Figs. 5 and 6). Thus, the precisions in mass improve only marginally when matching all three observables against stellar models and those in age improve significantly only when the third observable adds valuable new information, e.g. if the density of isochrones is high (close to the ZAMS).

The change of the precisions when incorporating correlations is a combination of the changes discussed in Secs. 3.2.1 and 3.2.2. For example, the mass precision is lower for stars in the HR diagram (Fig. 5b) and improves for stars in the Kiel diagram (Fig. 6b). In stars more massive than about $7 M_{\odot}$, the improvement due to correlations in effective temperature and gravity outweigh the lowering because of correlations in effective temperature and luminosity, while this is the opposite in stars less massive than $7 M_{\odot}$ (Fig. 7b). A similar situation is found for the age precisions (Fig. 7d): correlations in effective temperatures and luminosities improve the age precision for stars close to the ZAMS, while age precisions are lower for stars that are farther away (Fig. 5d). The opposite trend is found for correlations in effective temperatures and surface gravities (Fig. 6d).

Altogether the effects on the age precisions tend to compensate, but the correlation in effective temperature and gravity dominates and therefore outweighs the changes in the precisions.

3.3. Most likely inferred initial masses and ages

Not only the precision, but also the most likely inferred stellar parameters are modified when accounting for correlations. To judge the importance of these changes, we have to compare them to the inferred error bars. To this end, we systematically map the relative differences $(x_{\text{corr}} - x_{\text{no-corr}}) / \sigma_{\text{no-corr}}$ of inferred initial masses and ages in Fig. 8 (the indices “corr” and “no-corr” refer to the case when taking into account and neglecting correlations, respectively, and σ is the 1σ uncertainty of parameter x). As before, we consider the three cases separately when the positions of stars in the HR diagram, Kiel diagram, and both diagrams are known with 1σ uncertainties of 1000 K, 0.1 dex, and 0.1 dex in effective temperature, luminosity, and surface gravity, respectively. In all cases the differences are within the 1σ uncertainties of the stellar parameters, but still result in systematic

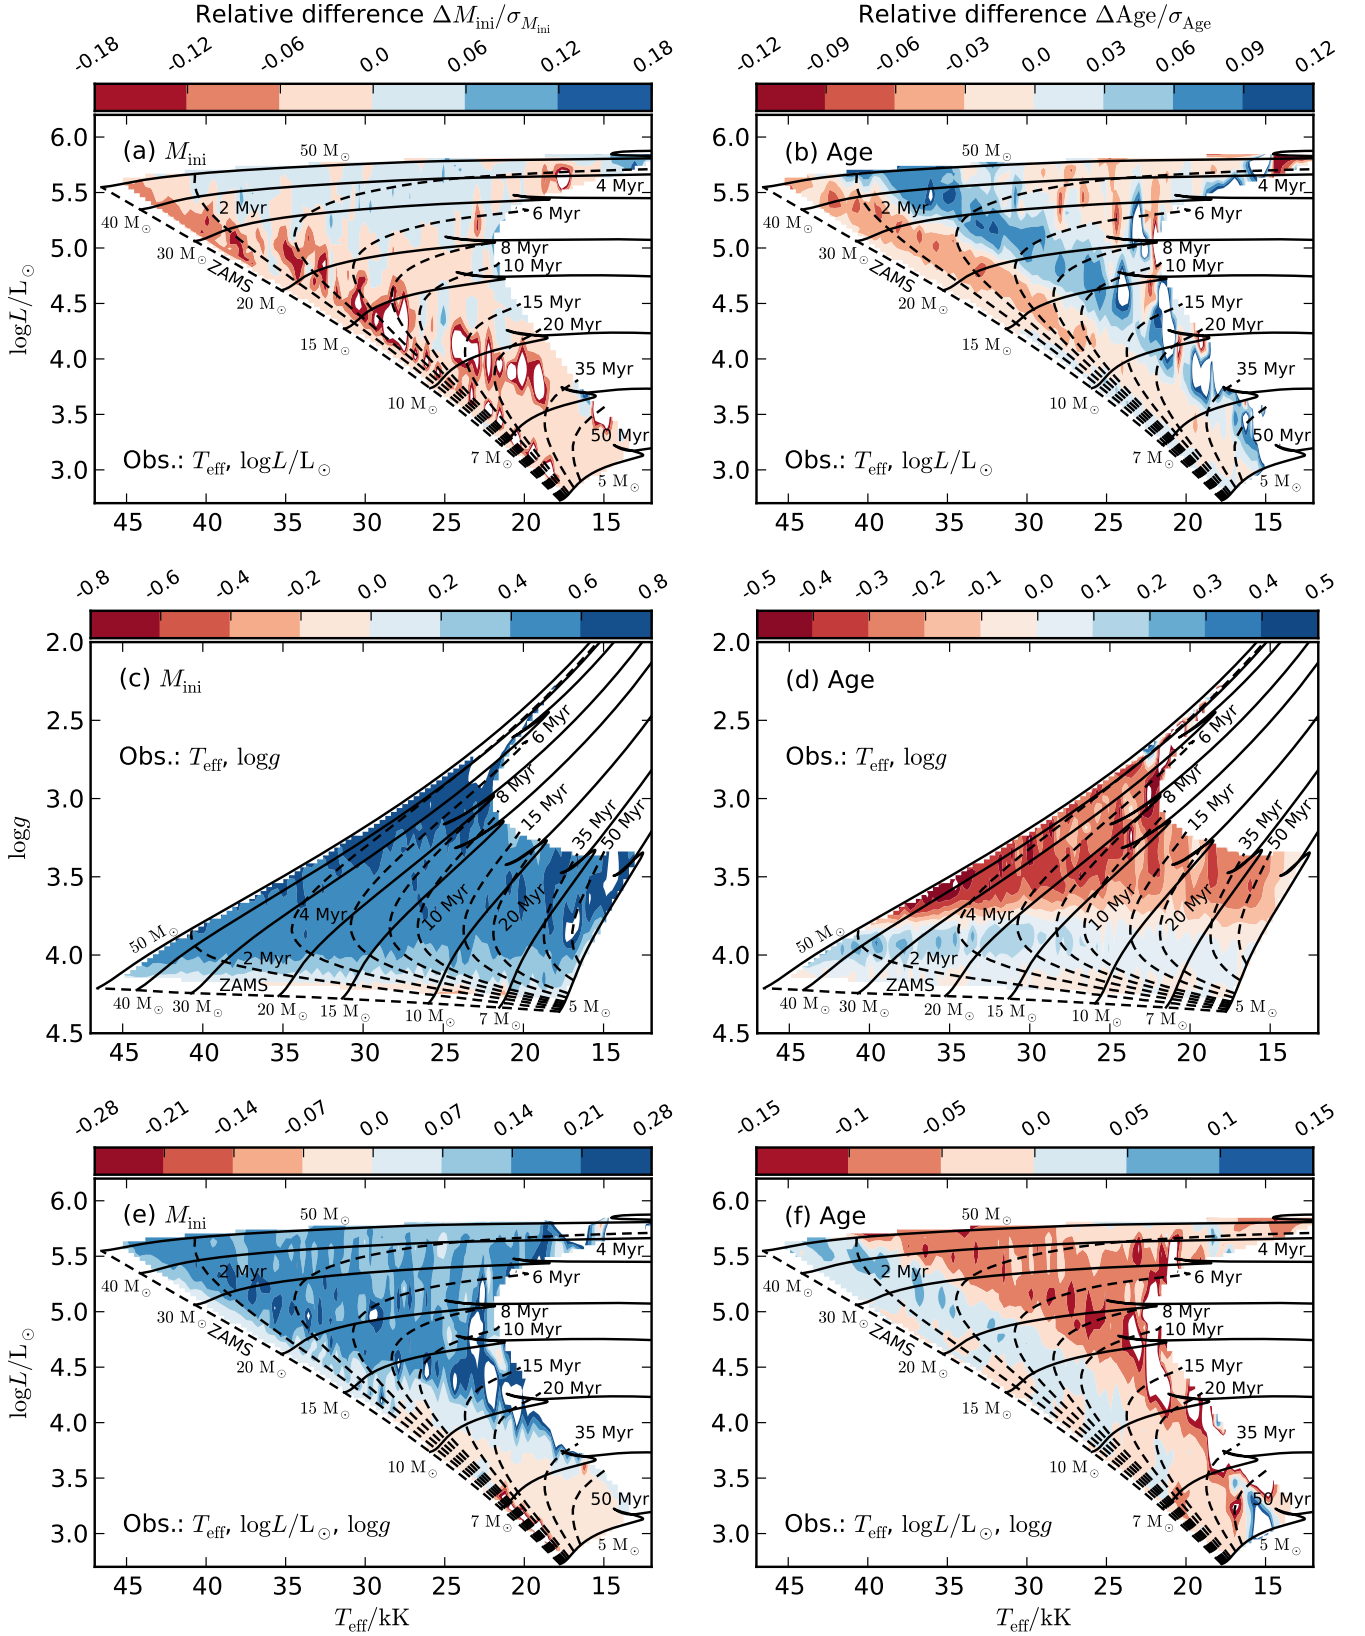


Fig. 8. Relative differences in the inferred most likely initial mass (left panels) and age (right panels) when taking correlations into account. In the top panels (a) and (b) we show the relative differences for stars with known effective temperatures and luminosities, in the middle panels (c) and (d) for stars with known effective temperatures and surface gravities, and in the bottom panels (e) and (f) for stars with known effective temperatures, luminosities, and surface gravities. The differences are defined with respect to the most likely parameters when neglecting correlations, i.e. $(x_{\text{corr}} - x_{\text{no-corr}})/\sigma_{\text{no-corr}}$ where x is the parameter and σ the 1σ uncertainty. The indices “corr” and “no-corr” show that correlations have been taken into account and neglected, respectively. The different ranges in the relative differences in each panel should be noted.

biases when neglecting correlations and may thus be particularly important when considering samples of stars.

As is evident from Figs. 8a and 8b, the differences in initial mass and age for our stars in the HR diagram are of the order of $\pm 5\%$ of the 1σ error bars and at most about 20% and 15%, respectively. The correlations are stronger for our stars in the Kiel diagram, leading to larger relative differences in the most likely inferred masses and ages (Figs. 8c and 8d). Masses are systematically more massive by about 50% of the 1σ error bars. Ages can be systematically younger or older depending on the relative orientation of the error ellipse and isochrones. On average, ages change by about +10% and -25% of 1σ uncertainties. In the worst case, masses are underestimated by 0.8σ and ages over- and underestimated by 0.5σ and 0.3σ , respectively, showing the importance of correlations in this case. Matching the positions of stars in the HR and Kiel diagram simultaneously to stellar models results in slightly larger changes than what was found for stars in the HR diagram and smaller changes than for stars in the Kiel diagram. The differences are of the order of 15% in mass and 8% in age with respect to their corresponding 1σ uncertainties (Figs. 8e and 8f).

4. Conclusion

In this paper, we include the covariance matrix in the likelihood model of the Bayesian code BONNSAI (Schneider et al. 2014) to facilitate the use of correlated observables when matching observations of stars against stellar models. Because correlations are typically not published and often not available, we derive a parametrisation of the covariance matrix that requires conventional observables including their uncertainties and additionally a correlation parameter that describes how two observables covary and that only depends on the method used to determine the observables. The advantage of this parametrisation is that correlations can be easily incorporated and may even be used when information on correlations are not published with the data but are known in general.

Correlations modify the likelihood function and, in our case, the likelihood function has the shape of a rotated ellipse. Because of the rotation and change in the shape of the error ellipses, the likelihood function covers different parts of the parameter space of the observables, e.g. in the HR diagram, compared to the case when correlations are neglected. As a result of neglecting correlations, the most likely value of inferred stellar parameters such as mass and age can be systematically biased and the inferred error bars, i.e. the precision with which stellar parameters can be determined, can be under- or overestimated. We show that the importance of correlations depends on the interplay between the strength of correlations and the conventional error bars.

In our example of OB stars with effective temperatures and surface gravities being determined from atmosphere modelling, the correlations are significant and we find that the precision with which masses can be derived improves by about a factor of 2. At the same time the precision in age decreases over a wide range of effective temperatures and surface gravities. We also find that initial masses are systematically underestimated on average by 0.5σ and ages often systematically overestimated when neglecting correlations for these stars in the Kiel diagram. In all of the cases, the reason for the differences is the change in the orientation of the likelihood model with respect to the stellar tracks and isochrones.

A likelihood model that takes correlations properly into account is a better approximation of the data. The reliability and

robustness of inferred fundamental stellar parameters is therefore always enhanced, manifesting itself in systematic changes in the precision and most likely values of inferred stellar parameters. This fosters the importance of taking correlations properly into account when approaching an era of precision stellar astrophysics with current and upcoming surveys such as Gaia, and whenever robust error bars are essential.

Acknowledgements. We thank J. Puls for helpful discussions regarding correlations from stellar atmosphere modelling, and an anonymous referee for carefully reading our manuscript and making constructive suggestions. FRNS thanks the Hintze Family Charitable Foundation and Christ Church College for his Hintze and postdoctoral research fellowships.

Appendix A: Influence of prior distributions

In a Bayesian analysis, not only does the likelihood model influence inferred model parameters (precisions and most likely values), but also the prior functions. In our default statistical model, we use a Salpeter IMF as initial mass prior distribution, a Gaussian as initial rotational velocity prior distribution, and a uniform age prior distribution, which are adequate descriptions of prior knowledge of massive, Milky Way stars.

A Salpeter IMF prior distribution puts more weight on lower mass stars. Any M_{ini} -prior distribution influences the inference of initial masses most strongly in a region of the observational parameter space where tracks of different initial mass are closest to each other. For example, in the HR diagram stellar tracks are closer to each other when the initial mass is higher. Changing the M_{ini} -prior distribution will therefore have a greater influence on higher mass stars than on lower mass stars (assuming constant error bars). The main influence of the M_{ini} -prior distribution is to change the most likely values of inferred initial masses and ages; for example, replacing the Salpeter IMF prior by a uniform M_{ini} -prior distribution for stars in the HR diagram changes the precision with which mass and age can be determined by up to $\pm 5\%$, and shifts all masses to higher values and ages to lower values because of the anti-correlation of mass and age (more massive stars have shorter lifetimes). The shift is strongest for the most massive stars and can be up to $+0.4\sigma$ in mass and -0.4σ in age (Fig. A.1).

Rotation can significantly change stellar models. In general, rotation “blurs” tracks of the same initial mass in the HR diagram such that a wider variety of initial masses can reach the same effective temperatures and luminosities. Giving more weight to a broader range of initial rotational velocities, for example by choosing a greater width of the Gaussian v_{ini} -prior distribution or by even using a uniform prior distribution, a wider range of models becomes likely to explain the same position of stars in the HR diagram, hence lowering the precisions with which mass and age can be inferred. Furthermore, rotating stars can be either more or less luminous than non-rotating stars depending on their mass and evolutionary stage. This is due to a balance of rotationally induced chemical mixing making stars more luminous and centrifugal forces making them less luminous (on the ZAMS, rotating stellar models are therefore always less luminous because mixing has had no effect yet). In the Milky Way models of Brott et al. (2011), (rapidly) rotating stars more massive than about $20 M_{\odot}$ can become considerably more luminous than their non-rotating counterparts during the MS evolution. Giving more weight to rapid rotators by changing (broadening) the v_{ini} -prior distribution will therefore shift the most likely mass to smaller values for $M_{\text{ini}} \gtrsim 20 M_{\odot}$ because less massive, rapid rotators reaching the same effective temperatures and luminosities as slowly rotating, more massive

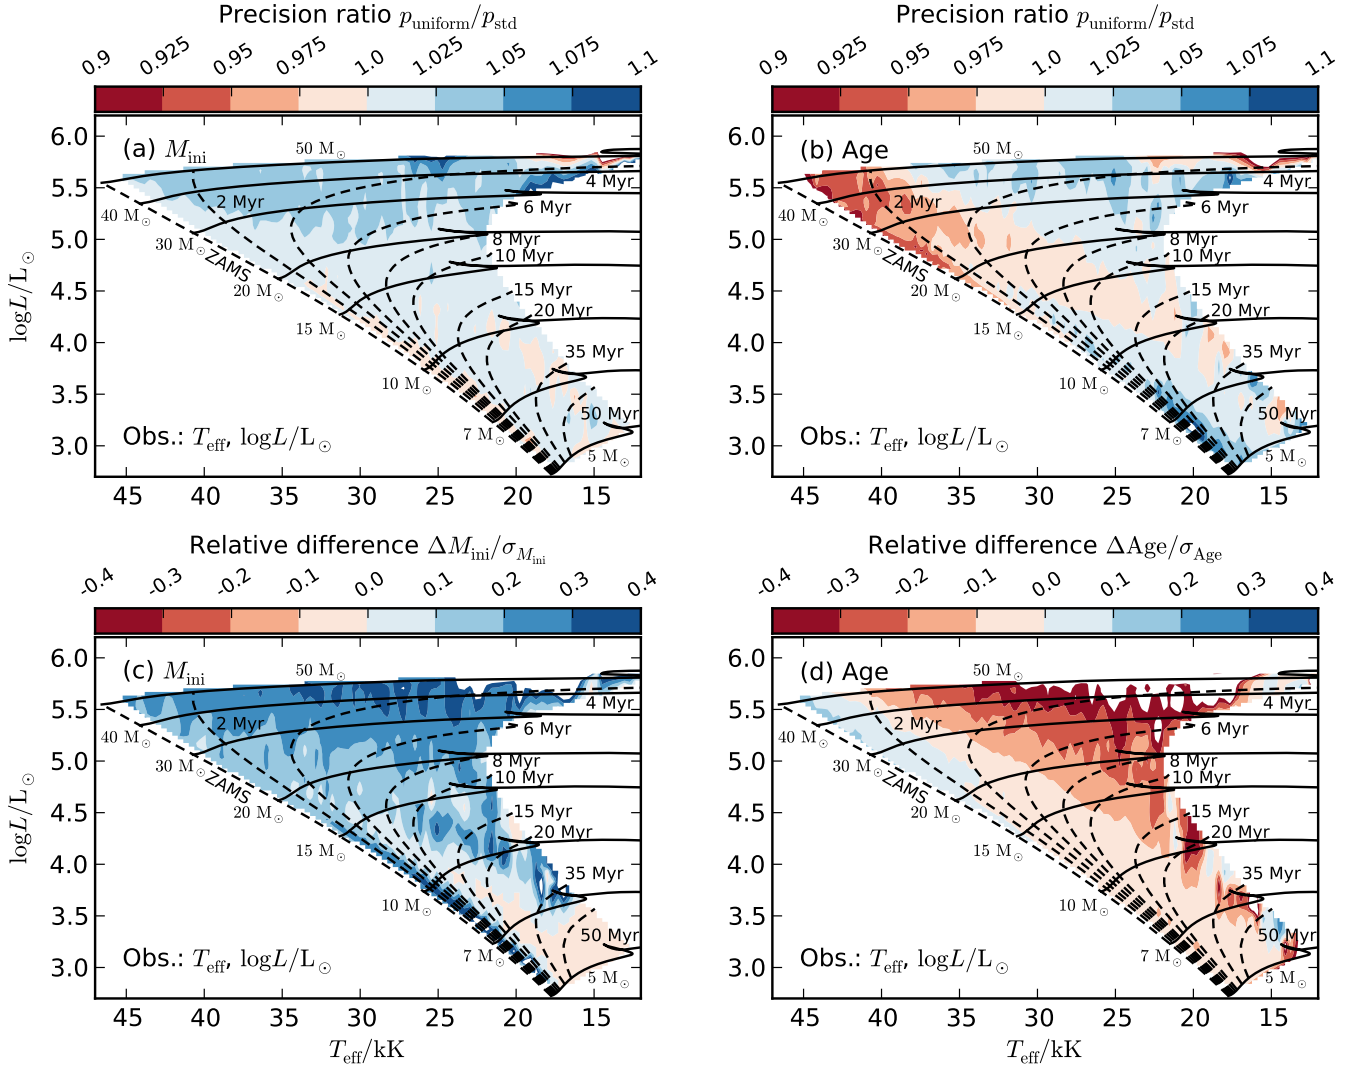


Fig. A.1. Changes in the precisions $p = \sigma_x/x$ (top row) and most likely values (bottom row) of initial mass (left column) and age (right column) when assuming a uniform M_{ini} -prior instead of the Salpeter IMF prior distribution applied in our default statistical model (the changes in most likely values are defined analogously to those in Fig. 8, i.e. $\Delta x = x_{\text{uniform}} - x_{\text{std}}$; x_{uniform} and x_{std} , and p_{uniform} and p_{std} denote the most likely values and precisions for a uniform M_{ini} -prior distribution and our standard prior choice, respectively). The observables are effective temperature and luminosity, and the precision ratios and relative differences are with respect to the precisions and most likely values when neglecting correlations (cf. Figs. 5, 8a, and 8b). For more details, see Sec. A.

stars are no longer disfavoured as much as before by the prior distribution (or are not disfavoured at all for a uniform prior distribution). Analogously, the most likely masses shift to larger values for $M_{\text{ini}} \lesssim 20 M_{\odot}$, and the most likely ages will be older for $M_{\text{ini}} \gtrsim 20 M_{\odot}$ and younger for $M_{\text{ini}} \lesssim 20 M_{\odot}$. Close to the ZAMS, inferred masses always shift to larger values and ages to lower values. In the most extreme (and unrealistic) case of using a uniform v_{ini} -prior distribution, the precisions of inferred masses and ages is reduced by up to 20% and 60%, respectively. At the same time, the most likely masses and ages can change by up to ± 0.5 – 0.6σ (Fig. A.2).

The changes in precision and most likely mass and age quoted above can be regarded as upper limits because they originate from changing either the M_{ini} - or the v_{ini} -prior distribution by a maximum extent such that all initial masses or initial rotational velocities are a priori equally probable. This situation corresponds to going from a simple maximum-likelihood approach (uniform prior distributions) to a full Bayesian approach. The

influence of changing prior distributions on inferred model parameters are qualitatively the same for stars in the Kiel diagram, but differ quantitatively.

References

- Bonfanti, A., Ortolani, S., & Nascimbeni, V. 2016, *A&A*, 585, A5
- Brott, I., de Mink, S. E., Cantiello, M., et al. 2011, *A&A*, 530, A115
- Burnett, B. & Binney, J. 2010, *MNRAS*, 407, 339
- Castro, N., Urbaneja, M. A., Herrero, A., et al. 2012, *A&A*, 542, A79
- da Silva, L., Girardi, L., Pasquini, L., et al. 2006, *A&A*, 458, 609
- Flower, P. J. 1996, *ApJ*, 469, 355
- Fossati, L., Castro, N., Morel, T., et al. 2015, *A&A*, 574, A20
- Fossati, L., Schneider, F. R. N., Castro, N., et al. 2016, *A&A*, in press
- Hunter, I., Lennon, D. J., Dufton, P. L., et al. 2008, *A&A*, 479, 541
- Jones, S., Hirschi, R., Pignatari, M., et al. 2015, *MNRAS*, 447, 3115
- Jørgensen, B. R. & Lindegren, L. 2005, *A&A*, 436, 127
- Kudritzki, R.-P. & Puls, J. 2000, *ARA&A*, 38, 613
- Lastennet, E. & Valls-Gabaud, D. 2002, *A&A*, 396, 551
- Martins, F. & Palacios, A. 2013, *A&A*, 560, A16
- Macted, P. F. L., Serenelli, A. M., & Southworth, J. 2015a, *A&A*, 575, A36

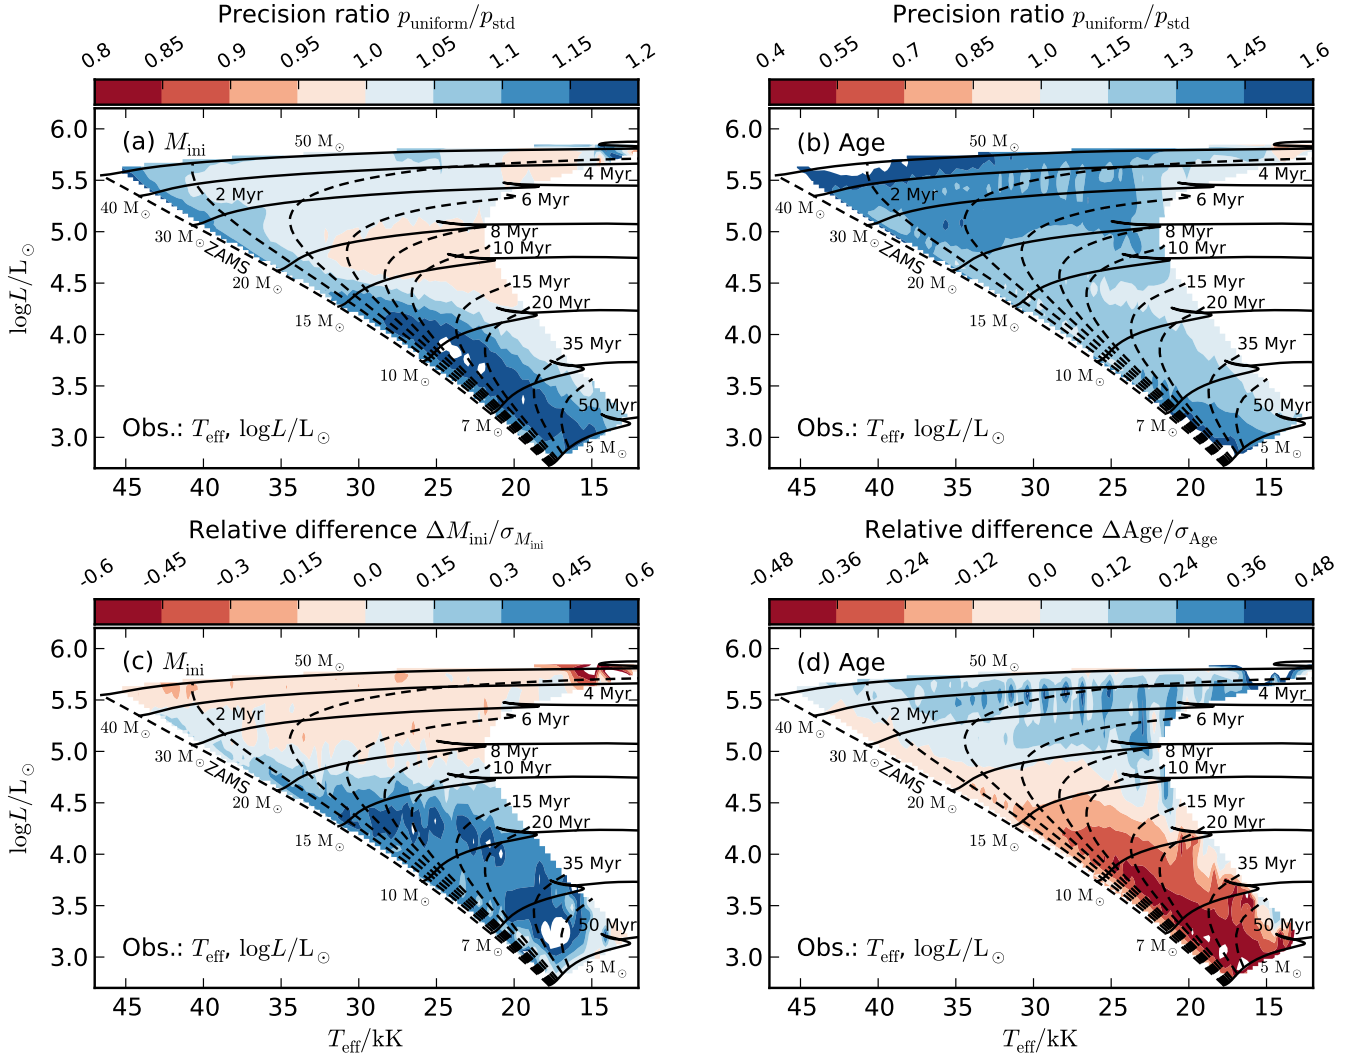


Fig. A.2. Same as Fig. A.1 but applying a uniform v_{ini} -prior instead of the Gaussian v_{ini} -prior distribution used in our default statistical model. As M_{ini} -prior distribution, we use a Salpeter IMF here. For more details, see Sec. A.

Maxted, P. F. L., Serenelli, A. M., & Southworth, J. 2015b, *A&A*, 577, A90
 Mitschang, A. W., De Silva, G., Zucker, D. B., et al. 2014, *MNRAS*, 438, 2753
 Pont, F. & Eyer, L. 2004, *MNRAS*, 351, 487
 Puls, J., Kudritzki, R.-P., Herrero, A., et al. 1996, *A&A*, 305, 171
 Puls, J., Urbaneja, M. A., Venero, R., et al. 2005, *A&A*, 435, 669
 Salpeter, E. E. 1955, *ApJ*, 121, 161
 Santolaya-Rey, A. E., Puls, J., & Herrero, A. 1997, *A&A*, 323, 488
 Schneider, F. R. N., Langer, N., de Koter, A., et al. 2014, *A&A*, 570, A66
 Schönnrich, R. & Bergemann, M. 2014, *MNRAS*, 443, 698
 Serenelli, A. M., Bergemann, M., Ruchti, G., & Casagrande, L. 2013, *MNRAS*, 429, 3645
 Shkedy, Z., Decin, L., Molenberghs, G., & Aerts, C. 2007, *MNRAS*, 377, 120
 Simón-Díaz, S., Castro, N., García, M., Herrero, A., & Markova, N. 2011, *Bulletin de la Société Royale des Sciences de Liege*, 80, 514
 Simón-Díaz, S., Negueruela, I., Mafé Apellániz, J., et al. 2015, in *Highlights of Spanish Astrophysics VIII*, ed. A. J. Cenarro, F. Figueras, C. Hernández-Monteagudo, J. Trujillo Bueno, & L. Valdivielso, 576–581
 Stancliffe, R. J., Fossati, L., Passy, J.-C., & Schneider, F. R. N. 2015, *A&A*, 575, A117
 Stancliffe, R. J., Fossati, L., Passy, J.-C., & Schneider, F. R. N. 2016, *A&A*, 586, A119
 Takeda, G., Ford, E. B., Sills, A., et al. 2007, *ApJS*, 168, 297
 Torres, G., Andersen, J., & Giménez, A. 2010, *A&A Rev.*, 18, 67
 van Dyk, D. A., Degennaro, S., Stein, N., Jefferys, W. H., & von Hippel, T. 2009, *Annals of Applied Statistics*, 3, 117